

1 METHOD AN APPARATUS FOR DETERMINING JITTER BUFFER SIZE
2 IN A VOICE OVER PACKET COMMUNICATIONS SYSTEM
3

4 This application claims the benefit of co-owned, co-pending,
5 provisional application serial number 60/276,630 filed March 16,
6 2001, entitled "Methods and Apparatus for Delivering Multimedia
7 Communications Services to Multiple Tenants in a Building," the
8 complete disclosure of which is hereby incorporated by reference
9 herein.
10

11 BACKGROUND OF THE INVENTION
12

13 1. Field of the Invention

14 The invention relates to packet switching networks. More
15 particularly, the invention relates to the transmission of voice
16 signals over a packet network.
17

18 2. State of the Art

19 Before the advent of digital communications, voice telephone
20 communications was accomplished by setting up a discrete physical
21 circuit between the caller and the called party.
22

23 The first commercial digital voice communications system was
24 installed in 1962 in Chicago, Illinois. The system was called
25 "T1" and was based on the time division multiplexing (TDM) of

1 twenty-four telephone calls on two twisted wire pairs. The
2 digital bit rate of the T1 system was 1.544 Mbit/sec (± 200 bps),
3 which was, in the nineteen sixties, about the highest data rate
4 that could be supported by a twisted wire pair for a distance of
5 approximately one mile. The cables carrying the T1 signals were
6 buried underground and were accessible via manholes, which were,
7 at that time in Chicago, spaced approximately one mile apart.
8 Thus, analog amplifiers with digital repeaters were conveniently
9 located at intervals of approximately one mile.

10
11 The T1 system is still widely used today and forms a basic
12 building block for higher capacity communication systems including
13 T3 which transports twenty-eight T1 signals. The designation T1
14 was originally coined to describe a particular type of carrier
15 equipment. Today T1 is often used to refer to a carrier system, a
16 data rate, and various multiplexing and framing conventions.
17 While it is more accurate to use the designation "DS1" when
18 referring to the multiplexed digital signal carried by the T1
19 carrier, the designations DS1 and T1 are often used
20 interchangeably. Today, T1/DS1 systems still have a data rate of
21 1.544 Mbit/sec and support typically twenty-four voice and/or data
22 DS0 channels. Similarly, the designations DS2 and T2 both refer
23 to a system transporting up to four DS1 signals (96 DS0 channels)
24 and the designations DS3 and T3 both refer to a system
25 transporting up to seven DS2 signals (672 DS0 channels). The

1 timing tolerance for modern T1 equipment has been raised to ± 50
2 bps. T1 signals are said to be "plesiochronous" (nearly
3 synchronous). Clock variations at nodes are corrected by line
4 codes such as alternate mark inversion (AMI or bipolar line code).
5 These codes set up patterns in the bitstream of the signal which
6 are used at nodes to correct for clock variations.

7
8 Today, another higher bandwidth TDM system is in use. This
9 system is referred to as the synchronous optical network (SONET)
10 or, in Europe, the synchronous digital hierarchy (SDH). Unlike
11 plesiochronous signals, SONET signals are synchronized to a master
12 network clock. Although the timing of SONET signals is very
13 accurate, some clock variations still exist at different nodes in
14 the network. Various complex techniques are provided to correct
15 for clock differences at different nodes.

16
17 The T1 and T3 networks were originally designed for digital
18 voice communication. In a voice network minor bit errors can be
19 tolerated as a small amount of noise. However, in a data network,
20 a minor bit error cannot be tolerated. In the early 1970s,
21 another technology was deployed to support data networks. The
22 technology was called "packet switching". Unlike the T1 and T3
23 networks, packet switching was designed for data communications
24 only. In packet switching, a "packet" of data includes a header,
25 a payload, and a cyclic redundancy check (CRC). The header

1 includes addressing information as well as an indication of the
2 length of the payload. The payload contains the actual data which
3 is being transmitted over the network. The CRC is used for error
4 detection. The receiver of the packet performs a calculation with
5 the bits in the packet and compares the result of the calculation
6 to the CRC value. If the CRC value is not the same as the result
7 of the calculation, it means that the packet was damaged in
8 transit. According to the packet switching scheme, the damaged
9 packet is discarded and the receiver sends a message to the
10 transmitter to resend the packet. One popular packet switching
11 scheme for wide area networks (WANs), known as X.25, utilizes a
12 packet which has a fixed payload of 128 octets. Other packet
13 switching schemes allow variable length packets up to 2,000
14 octets. Frame Relay is an example of a WAN packet switching
15 scheme which utilizes variable sized packets and Ethernet is an
16 example of a local area network (LAN) packet switching scheme
17 which utilizes variable sized packets. Packet switching networks
18 are asynchronous and, by design, are not well suited for the
19 transmission of a streaming signal such as voice or video. If
20 streaming voice or video is transmitted via packets, small amounts
21 of noise in the signal will result in discontinuity of the stream,
22 echo, and other problems.

23
24 Concurrent with the development of packet switching several
25 groups around the world began to consider standards for the

1 interconnection of computer networks and coined the term
2 "internetworking". The leading pioneers in internetworking were
3 the founders of ARPANET (the Advanced Research Projects Network).
4 ARPA, a U.S. Department of Defense organization, developed and
5 implemented the transmission control protocol (TCP) and the
6 internet protocol (IP). The TCP/IP code was dedicated to the
7 public domain and was rapidly adopted by universities, private
8 companies, and research centers around the world. An important
9 feature of IP is that it allows fragmentation operations, i.e. the
10 segmentation of packets into smaller units. This is essential to
11 allow networks which utilize large packets to be coupled to
12 networks which utilize smaller packets. Today, TCP/IP is the
13 foundation of the Internet. It is used for email, file transfer,
14 and for browsing the Worldwide Web. It is so popular that many
15 organizations are hoping to make it the worldwide network standard
16 for all types of communication, including voice and video.

18 Perhaps the most awaited, and now fastest growing technology
19 in the field of telecommunications is known as Asynchronous
20 Transfer Mode (ATM) technology. ATM was originally conceived as a
21 carrier of integrated traffic, e.g. voice, data, and video. ATM
22 utilizes fixed length packets (called "cells") of 53 octets (5
23 octets header and 48 octets payload). ATM may be implemented in
24 either a LAN or a WAN. For ideal data transfer, ATM is used end
25 to end from the data source to the data receiver.

1 Current ATM service is offered in different categories
2 according to a user's needs. Some of these categories include
3 constant bit rate (CBR), variable bit rate (VBR), unspecified bit
4 rate (UBR), and available bit rate (ABR). CBR service is given a
5 high priority and is used for streaming data such as voice and
6 video where a loss of cells would cause a noticeable degradation
7 of the stream and where retransmission of lost cells is pointless.
8 UBR and ABR services are given a low priority and are used for
9 data transfers such as email, file transfer, and web browsing
10 where sudden loss of bandwidth (bursty bandwidth) can be tolerated
11 and lost cells can be retransmitted beneficially. ATM service is
12 sometimes referred to as "statistical multiplexing" as it attempts
13 to free up bandwidth which is not needed by an idle connection for
14 use by another, active, connection.

15
16 ATM switches (like other packet switches) typically include
17 multiple ("jitter") buffers, queues, or FIFOs for managing the
18 flow of ATM cells through the switch. Generally, a separate
19 buffer is provided for each outlet from the switch. However, it
20 is also known to have separate buffers at the inlets to the
21 switch. Buffer thresholds are set to prevent buffer overflow. If
22 the number of cells in a buffer exceeds the threshold, no more
23 cells are allowed to enter the buffer. Cells attempting to enter
24 a buffer which has reached its threshold will be discarded.

1 Whenever streaming real time data such as voice or video is
2 transmitted over a packet network, careful attention must be paid
3 to the size of the jitter buffers in the packet switches. If the
4 jitter buffer is too small, it can overflow resulting in the loss
5 of data. In a voice over packet system, buffer overflow is
6 usually perceived as an audible click or pop. If the jitter
7 buffer is too large, the data will be subject to unnecessary
8 delay.

9
10 There are two significant thresholds that are crossed as
11 network delay increases. The first of these thresholds results
12 from an electrical quirk of the world's telephone network.
13 Telephone handsets use 4 wires to connect to the network.
14 However, the network only uses a single pair of wires to carry a
15 telephone call from a telephone to a central office. This
16 transition from 2-wire to 4-wire operation is accomplished by the
17 use of a device known as a "hybrid" in the telephones at each end
18 of the network connection. Hybrids cause a reflection of signal,
19 due to an unavoidable impedance mismatch. The effect of this is
20 that some speech signal is reflected back toward the speaker by
21 each hybrid. Normally this reflection is not noticed. However,
22 if there is delay across the network, the reflected signal from
23 the far-side hybrid will arrive back as an echo. The effect of
24 this echo is more pronounced as the delay increases. At about 30
25 milliseconds of network delay, the echo is so significant as to

1 make normal conversation very difficult. Therefore, once a voice
2 circuit delay exceeds 30 milliseconds, echo cancelers must be
3 included within the network. These devices are expensive and
4 complex. Moreover, they operate most successfully when the delay
5 over a circuit is constant, and this may not be the case in a
6 packetized voice network.

7
8 Once in place echo cancellation systems allow network delays
9 to reach approximately 150 milliseconds before further voice
10 quality degradation is experienced. At about 150 milliseconds of
11 delay a second problem emerges. At this level of delay, most
12 people start to have significant problems in carrying on a normal
13 conversation. Normal conversation patterns demand that some
14 responses from the listener be received within less than 200
15 milliseconds, and delays of this order result in stilted
16 conversations, and "clashing" (where both parties try to talk at
17 once). This type of problem is commonly encountered over
18 satellite links.

19
20 For these reasons toll quality networks require that the end
21 to end network delay for voice traffic is less than 25
22 milliseconds in national networks and less than 100 milliseconds
23 in an international context.

1 Buffering is not the only source of delay in a packet
2 network. Delay also results from analog to digital conversion,
3 packetization, and digital to analog conversion. Nevertheless,
4 buffering must be given careful attention in order to maintain
5 acceptable voice quality.

6
7 One of the difficulties in choosing buffer size in a packet
8 network is that the jitter characteristics of the network vary
9 over time depending on network usage. Current methods for
10 determining the proper size of a jitter buffer consider the worst
11 case scenario to compensate for network jitter during peak hours
12 at the expense of introducing extra delay during non-peak hours.

13
14 The present Internet standard for transporting realtime
15 streaming data such as voice and video is contained in RFC-1889
16 which was published in January 1996, the complete disclosure of
17 which is hereby incorporated by reference herein. RFC-1889
18 describes the "real-time transport protocol" (RTP), which provides
19 end-to-end delivery services for data with real-time
20 characteristics, such as interactive audio and video. Those
21 services include payload type identification, sequence numbering,
22 timestamping, and delivery monitoring. Applications typically run
23 RTP on top of UDP (user datagram protocol) to make use of UDP's
24 multiplexing and checksum services. Both protocols contribute to
25 the transport protocol functionality. However, RTP may be used

1 with other suitable underlying network or transport protocols.
2 RTP supports data transfer to multiple destinations using
3 multicast distribution if available in the underlying network.
4

5 RTP itself does not provide any mechanism to ensure timely
6 delivery of packets or provide any other quality-of-service
7 guarantees, but relies on lower-layer services to do so. It does
8 not guarantee delivery or prevent out-of-order delivery, nor does
9 it assume that the underlying network is reliable and delivers
10 packets in sequence. The sequence numbers included in RTP allow
11 the receiver to reconstruct the sender's packet sequence, but
12 sequence numbers might also be used to determine the proper
13 location of a packet, for example in video decoding, without
14 necessarily decoding packets in sequence.
15

16 SUMMARY OF THE INVENTION 17

18 It is therefore an object of the invention to provide methods
19 and apparatus for determining the proper size for jitter buffers
20 in a voice over packet network.
21

22 It is also an object of the invention to provide methods and
23 apparatus for determining the proper size for jitter buffers in a
24 voice over packet network where the jitter characteristic of the
25 network varies with time.

1 It is another object of the invention to provide methods and
2 apparatus for determining the proper size for jitter buffers which
3 does not allow undue delay during off-peak hours.
4

5 It is still another object of the invention to provide
6 methods and apparatus for determining the proper size for jitter
7 buffers such that the size of the buffers can be adjusted based on
8 a quantitative measurement of actual network jitter.
9

10 In accord with these objects which will be discussed in
11 detail below, the methods of the present invention are applicable
12 to a packet switch which sends and receives RTP packets, has a
13 local clock and an adjustable dynamic buffer, and where the size
14 of the packets received is relatively constant. The methods of
15 the invention include determining packet size by comparing the RTP
16 timestamps of two consecutive packets, determining the expected
17 (no jitter) local arrival time for the next packet by comparing
18 the difference between the local clock and the timestamp of the
19 present packet and summing it with the timestamp difference of the
20 last two packets, and determining network jitter by comparing the
21 expected local arrival time with the actual local arrival time.
22 The absolute values of computed network jitter are averaged to
23 determine the average network jitter. The average network jitter
24 is used to set the size of the dynamic buffer. According to the
25 presently preferred embodiment, large changes in jitter between

1 consecutive packets are considered anomalies and are not used in
2 the average jitter calculation. Conversely, only significant
3 changes in the average jitter will call for a readjustment of the
4 dynamic jitter buffer. When the buffer size is adjusted, audio
5 traffic experiences an audible click or pop. However, the
6 occasional click or pop is likely to be less annoying than undue
7 delay. The methods of the invention can also be used in
8 combination with a voice activity detector so that the buffer size
9 adjustments are made less noticeable. That is, the buffer size is
10 adjusted during a period of no voice activity.

11
12 Additional objects and advantages of the invention will
13 become apparent to those skilled in the art upon reference to the
14 detailed description taken in conjunction with the provided
15 figures.

16
17 BRIEF DESCRIPTION OF THE DRAWINGS

18
19 Figure 1 is a simplified block diagram of an apparatus
20 according to the invention; and

21
22 Figure 2 is a simplified flow chart illustrating the methods
23 of the invention.

1 DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

2
3 Referring now to Figure 1, an apparatus 10 according to the
4 invention is shown in conjunction with an RTP packet device 1
5 having a local clock 3, an adjustable dynamic buffer 5, and a
6 packet I/O 7. The apparatus 10 includes a packet sniffer 12
7 coupled to the packet I/O 7 for reading the RTP headers of
8 incoming packets, a logic block 14 for performing the methods of
9 the invention, and a group of registers 16 for storing values of
10 variables. The logic block 14 takes input from the packet sniffer
11 12 and the local clock 3, communicates bidirectionally with the
12 registers 16, and provides an adjustment output to the dynamic
13 buffer 5.
14

15 The invention utilizes the thirty-two bit timestamp which
16 occupies bytes 4-7 of the RTP packet header. The timestamp
17 reflects the time of the first sample in the packet payload based
18 on the local clock of the transmitting device. The sample period
19 for G.711 A-law/ μ -law is always 125 μ s (8000hz). When the RTP
20 device 1 is powered up, the timestamp is set to a random thirty-
21 two bit value and is incremented by the local clock 3 until the
22 device is shut down, wrapping to zero as the value overflows. The
23 prior art devices only use the local clock to generate the
24 timestamp in outgoing packets. The present invention uses the

1 local clock and the timestamp on incoming packets to determine
2 network jitter.

3
4 Turning now to Figure 2, the methods of the invention are
5 illustrated in a flow chart starting at 100. The first step in
6 the method is to determine the length of the decompressed voice
7 signal. It is not sufficient to simply look at the packet size
8 because different voice coders will produce different packet sizes
9 for the same amount of time. For example, ADPCM coding compresses
10 1 ms to 4 bytes, whereas μ -law coding compresses 1 ms to 8 bytes.
11 By calculating the RTP timestamp difference (ΔR) between two
12 consecutive packets, the expanded packet size (length of the
13 decompressed voice signal) can be determined. No matter what
14 compression system is used, the RTP timestamps are always based on
15 an 8 KHz clock. Therefore, the difference ΔR is the duration in
16 125 μ s periods. Thus, the first step shown in Figure 2 is to
17 calculate ΔR as indicated at 102. This value is stored in one of
18 the registers (16 in Figure 1). In a zero jitter network, the
19 expected local arrival time of the next packet will be the local
20 arrival time of the last packet plus ΔR . In order to calculate
21 the expected local arrival time of the next packet in a zero
22 jitter network, the local arrival time (TL) of the last packet is
23 determined at 104 in Figure 2. This value is also stored in one
24 of the registers (16 in Figure 1). The calculation of the
25 expected local arrival time TE is illustrated at 106. The value

1 TE is also stored in one of the registers (16 in Figure 1). When
2 the next packet arrives, it will arrive at an actual local time
3 TL' which, in a network with jitter, will be different from the
4 expected time TE. The actual arrival time of the next packet is
5 determined at 108 and this value TL' is stored in one of the
6 registers (16 in Figure 1). The difference between the actual
7 arrival time TL' and the expected arrival time TE is calculated at
8 110 and the difference ΔN is considered an indication of network
9 jitter. The absolute value of ΔN is stored in another one of the
10 registers (preferably an accumulator). A count of the number of
11 ΔN s accumulated is stored in another register. The average value
12 of the accumulated ΔN s is calculated at 112 and the result J is
13 stored in a register. Before proceeding to set the initial buffer
14 size, steps 102-112 are repeated for several hundred packets. For
15 example, a loop counter at 114 causes the calculation of J to be
16 updated until it is based on an average of several hundred packet
17 arrivals.

18
19 The buffer size is calculated at 116. According to the
20 presently preferred embodiment, the buffer size BS is calculated
21 according to the formula $BS = C + (S/2) + 2J$ where C is the minimum size
22 of a buffer in a zero jitter network, S is the size of the
23 packets, and J is the average network jitter as described above.
24 The buffer size BS is initially set at 118 and the method
25 continues at 102 through 112 to recalculate J based on the jitter

1 detected in the next received packet. According to the presently
 2 preferred embodiment, if the calculation at 110 produces a ΔN
 3 which differs greatly (e.g. by a packet size or more) from the
 4 last calculated ΔN , the calculation is ignored. After the new
 5 buffer size BS' is calculated at 116, it is compared at 120 to the
 6 present buffer size BS. If the difference is insignificant (e.g.
 7 less than 8 ms), no change is made in the buffer size and the
 8 method returns to 102. If the difference is substantial (e.g.
 9 more than 8 ms), the buffer size is adjusted at 122 and the
 10 process proceeds at 102 through continuing cycles.

11
 12 In the presently preferred implementation, the methods of the
 13 invention are separated into a foreground task and a background
 14 task. The foreground task calculates J and the background task
 15 updates the buffer size. Thus, in the Flow chart, the steps 118,
 16 120, and 122 are run in a separate timed loop. Exemplary source
 17 code for the foreground task is illustrated in the code listing
 18 below which is separately line numbered for easy reference.

```

1      ; Detect Jitter
2      ;
3
4      ; don't detect jitter until channel has reached a steady state
5      condition
6
7      rsbx SXM
8      ldm    T, B
9      stlm   B, AR4
10     mpy     #HEADER_SIZE, B
11     add     #_chnl_header, B
12     stlm    B, AR2
  
```



```

13      ld          *AR4(_chnl_state), B
14      and          #RX_ON_MASK, B          ; Once spigot is
15      turned on, channel is in OK state
16      bc          UpdNPTime, BEQ
17
18      ; Channel is (hopefully) in a steady State condition
19      ; Compare Actual & Expected Arrival Times
20      dld          *AR3(_LastPacketLTime), B
21      ; Get Local Sample Time
22
23      ; load recieved (remote) rtp timestamp into acc. A.
24      ; This must be done as two loads because it may not be aligned
25      for dld
26      ld *AR2(LOCALSTAMP-1), 16, A      ; load high word
27      or *AR2(LOCALSTAMP), A          ; load low word
28
29      ;
30      ; B = predicted time time, A = actual time
31      ;
32      ssbx SXM
33      nop
34      ;xc          2, TC
35      ;stm          #0ffh, AG          ; correct AG for sign bit.
36      sub          A, B
37      abs          B
38      sub          #1000, B, A          ; If we have a huge difference, ignore it.
39      bc          UpdNPTime, AGT
40      dadd          *AR3(_AvgTimeDiff), B          ; add to the time diff's
41      dst          B, *AR3(_AvgTimeDiff)
42
43      UpdNPTime:
44
45      ;Update "Next Packet" Expected Arrival Time
46      mvdk          *AR1(UDP_HEADER_SIZE+RTP_TIMESTAMP_OFFSET-1), BH
47      mvdk          *AR1(UDP_HEADER_SIZE+RTP_TIMESTAMP_OFFSET), BL
48      ld          B, A
49      dsub          *AR3(_LastPacketRTime), B
50      st          #0, *(BG)
51      dst          A, *AR3(_LastPacketRTime)
52      mvdk
53      *AR2(ETHER_HEADER_SIZE+IP_HEADER_SIZE+UDP_HEADER_SIZE+RTP_TIMESTAMP_OFFSET-
54      1), AH
55      mvdk
56      *AR2(ETHER_HEADER_SIZE+IP_HEADER_SIZE+UDP_HEADER_SIZE+RTP_TIMESTAMP_OFFSET), AL
57      nop
58      add          A, B
59      dst          B, *AR3(_LastPacketLTime)
60

```

Lines 1-18 of the code listing determine whether the channel is in a steady enough state to begin calculating jitter. At lines 19 and 20, the actual and expected arrival times are compared. The local sample time is obtained at lines 21-44. The next packet expected arrival time is then updated at lines 46-60 and the process is repeated as described above with reference to Figure 2.

Exemplary source code for the foreground task is illustrated in the code listing below which is separately line numbered for easy reference.

```

1      case RING_CADENCE:
2      RingCadence++;
3      if (RingCadence >= 3)
4      RingCadence = 0;
5      writeFlag = 1;
6      msg[0] = HEARTBEAT;
7      msg[1] = 0xFF;
8      /* Use this opportunity to adjust jitter
9      buffers for optimum quality */
9      /* AvgTimeDiff is sum of absolute value
10     of *sample* time diffs */
11     for (line = 0; line <
12     NUMBER_OF_CHANNELS; line++)
13     if ((OctetsSent[line] > 32000)
14     && ((chnl_state[line] >> RX_SHIFT) & STATE_MASK) == ACTIVE_STATE)
15     {
16     if (CrashCount[line] >
17     0x100)
18     CrashCount[line]
19     = 0x100;
20     temp =
21     (AvgTimeDiff[line] / PacketsSent[line]) + (DataRate[line] >> 3) + 2 +
22     (AvgPcktRecvd[line] >> 1) + (CrashCount[line] << 3);
23     if (temp > 0x180)
24     temp = 0x180;
25     if (abs(temp -
26     OptimalJBLevel[line]) > 32)
27     OptimalJBLevel[line] = temp;
28
29

```

```

30         if (CrashCount[line] > 0)
31
32     CrashCount[line]--;
33     }
34     break;

```

1 The cycle time for resetting the buffer size is set at lines
 2 1-7 using the variables RING_CADENCE and HEARTBEAT. The average
 3 time difference is determined at lines 8-20. A possible new
 4 buffer-size is calculated at lines 21-25. At lines 26-27, it is
 5 determined whether the difference between the old buffer size and
 6 the new buffer size is more than 32 words. If it is, the new
 7 buffer size is used to adjust the buffer. CrashCount is used to
 8 increase the minimum buffer size every time the buffer crashes so
 9 as to attempt to prevent another crash.

11 There have been described and illustrated herein methods and
 12 apparatus for determining jitter buffer size and adjusting a
 13 dynamic buffer accordingly in a voice over packet communications
 14 system. While particular embodiments of the invention have been
 15 described, it is not intended that the invention be limited
 16 thereto, as it is intended that the invention be as broad in scope
 17 as the art will allow and that the specification be read likewise.
 18 It will therefore be appreciated by those skilled in the art that
 19 yet other modifications could be made to the provided invention
 20 without deviating from its spirit and scope as so claimed.

21